

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BOARD OF PATENT APPEALS AND INTERFERENCES**

In Re Application of:)	
)	
Nitzan Peleg, et al.)	Confirmation No. 6513
)	
Serial No.:10/691,175)	Examiner: Black, Linh
)	Group Art Unit: 2169
Filed: October 22, 2003)	
)	
For: Method and Apparatus for Performing)	HP Docket: 200308559-1
Conflict Resolution in Database)	TKHR Docket: 050849-1190
Logging)	

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an appeal from the final Office Action mailed October 7, 2008, rejecting claims 1-21 of the present application.

I. REAL PARTY IN INTEREST

The real party in interest of the instant application is Hewlett-Packard Development Company, a Texas Limited Liability Partnership having its principal place of business in Houston, Texas.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

III. STATUS OF THE CLAIMS

Claims 1-21 are pending in this application. All claims were rejected by the final Office Action and are the subject of this appeal.

IV. STATUS OF AMENDMENTS

There have been no claim amendments made after the final Office Action, and all amendments made before the final Office Action have been entered. The claim listing in section VIII. CLAIMS – APPENDIX (below) represents the present state of the claims.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

Embodiments of the claimed subject matter are summarized below with reference numbers and references to the written description (“specification”) and drawings. The subject matter described below appears in the original disclosure at least where indicated, and may further appear in other places within the original disclosure.

Embodiments according to independent claim 1 involve a system. The system comprises: a processor (p. 5 lines 18-24); a memory (p. 5 lines 18-24); a materialized view stored on the memory, the materialized view being derived at least in part from a table (FIG. 3 elements 128 and 130; p. 7 lines 1-5; p. 7 lines 15-24; p. 8 lines 4-14; p. 9 lines 10-22; p. 10 lines 4-6; p. 14 lines 5-10); a logging mechanism stored on the memory (FIG. 3 element 124), the logging mechanism configured to maintain a refresh log (p. 7 lines 1-5; p. 7 lines 15-22; p. 8 lines 5-15; p. 14 lines 5-15), the refresh log containing a first range and a second range (p. 10 lines 20-27; p. 14 lines 15-25) that at least partially overlap (p. 15 lines 1-10), the first range and the second range each having a timestamp associated therewith (p. 8 lines 10-15; p. 11 lines 1-10; p. 15 line 20 to p. 16 line 18), wherein the time stamp associated with each of the first range and second range respectively indicates when an operation corresponding to the first range and the second range occurred to the table (p. 8 lines 10-15); and a refresh manager stored on the memory (FIG. 3 element 126), the refresh manager configured to resolve conflicts (p. 14 lines 10-25; p. 15 lines 20-25) between the first range and the second range that at least partially overlap by selecting portions of the first range and the second range that have the more

recent timestamp and applying the selected portions of the first range and the second range to the materialized view (p. 16 lines 20 to p. 17 line 10; p. 17 lines 10-25; p. 18 lines 1-12).

Embodiments according to independent claim 5 involve a system. The system comprises: a processor (p. 5 lines 18-24); a memory (p. 5 lines 18-24); a materialized view stored on the memory, the materialized view being derived at least in part from a table (FIG. 3 elements 128 and 130; p. 7 lines 1-5; p. 7 lines 15-24; p. 8 lines 4-14; p. 9 lines 10-22; p. 10 lines 4-6; p. 14 lines 5-8); a logging mechanism stored on the memory (FIG. 3 element 124), the logging mechanism configured to maintain a refresh log (FIG. 2 element 100; FIG. 3 element 100; p. 7 lines 1-5; p. 7 lines 15-22; p. 8 lines 5-15; p. 14 lines 5-15), the refresh log containing a range and a single-row entry (p. 10 lines 20-27; p. 14 lines 15-25; p. 15 lines 10-20), the range and the single row entry each having a timestamp associated therewith (p. 8 lines 10-15; p. 11 lines 1-10; p. 15 line 20 to p. 16 line 18), wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table (p. 8 lines 10-15) and; a refresh manager stored on the memory (FIG. 3 element 126), the refresh manager configured to resolve conflicts (p. 14 lines 20-25; p. 15 lines 20-25) between the range and the single-row entry (p. 18 lines 10-20; p. 18 line 20 to p. 19 line by ignoring the single-row entry if the single-row entry is part of the range and if the single-row entry has the more recent timestamp and by applying the single-row entry to the materialized view if the single-row entry is not part of the range or if the range has the more recent timestamp (p. 19 lines 5-20; p. 20 lines 1-10; p. 20 lines 10-20; p. 21 lines 5-25).

Embodiments according to independent claim 10 involve a method. The method comprises: deriving a materialized view at least in part from a table (FIG. 3 element 130; p. 7 lines 1-5; p. 7 lines 15-24; p. 8 lines 4-14; p. 9 lines 10-22; p. 10 lines 4-6; p. 14 lines 5-8); ; storing a first range and a second range that at least partially overlap in a refresh log (FIG. 2

element 100; FIG. 3 element 100; p. 7 lines 1-5; p. 7 lines 15-22; p. 8 lines 5-15; p. 10 lines 20-27; p. 14 lines 15-25; p. 15 lines 1-10); associating a timestamp with the first range and the second range in the refresh log (p. 8 lines 10-15; p. 11 lines 1-10; p. 15 line 20 to p. 16 line 18) such that the time stamp associated with the first range indicates when an operation corresponding to the first range occurred to the table and the time stamp associated with the second range indicates when an operation corresponding to the second range occurred to the table (p. 8 lines 10-15); and resolving conflicts (p. 14 lines 20-25; p. 15 lines 20-25) between the first range and the second range by applying a portion of either the first range or the second range that has the more recent timestamp to the materialized view (p. 16 lines 20 to p. 17 line 10; p. 17 lines 10-25; p. 18 lines 1-12).

Embodiments according to independent claim 14 involve a method. The method comprises: deriving a materialized view at least in part from a table (FIG. 3 elements 128 and 130; p. 7 lines 1-5; p. 7 lines 15-24; p. 8 lines 4-14; p. 9 lines 10-22; p. 10 lines 4-6; p. 14 lines 5-8); storing a range and a single-row entry in a refresh log (FIG. 2 element 100; p. 7 lines 1-5; p. 7 lines 15-22; p. 8 lines 5-15; p. 10 lines 20-27; p. 14 lines 15-25), the range and the single-row entry each having a timestamp associated therewith (p. 8 lines 10-15; p. 11 lines 1-10; p. 15 line 20 to p. 16 line 18), wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry occurred to the table (p. 8 lines 10-15); ignoring the single-row entry if the single-row entry is part of the range and if the single-row entry has the more recent timestamp (p. 19 lines 5-20; p. 20 lines 1-10; p. 20 lines 10-20; p. 21 lines 5-25; and applying the single-row entry to the materialized view if the single-row entry is not part of the range or if the range has the more recent timestamp (p. 19 lines 5-20; p. 20 lines 1-10; p. 20 lines 10-20; p. 21 lines 5-25).

Embodiments according to independent claim 19 involve a computer program. The computer program comprises: a machine readable medium (p. 5 lines 18-24; p. 29 lines 20-25); a logging mechanism stored on the machine readable medium (FIG. 3 element 124), the logging mechanism being adapted to create a refresh log (p. 7 lines 1-5; p. 7 lines 15-22; p. 8 lines 5-15; p. 14 lines 5-15) that contains a first range and a second range (p. 10 lines 20-27; p. 14 lines 15-25) that at least partially overlap (p. 15 lines 1-10), the first range and the second range each having a timestamp associated therewith (p. 8 lines 10-15; p. 11 lines 1-10; p. 15 line 20 to p. 16 line 18), wherein the time stamp associated with each of the first range and second range respectively indicates when an operation corresponding to the first range and the second range occurred to the table (p. 8 lines 10-15); and a refresh manager (FIG. 3 element 126) stored on the machine readable medium, the refresh manager being adapted to resolve conflicts (p. 14 lines 20-25; p. 15 lines 20-25) between the first range and the second range that at least partially overlap by selecting portions of the first range and the second range that have the more recent timestamp and applying the selected portions of the first range and the second range to the materialized view (p. 16 lines 20 to p. 17 line 10; p. 17 lines 10-25; p. 18 lines 1-12).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The following grounds of rejection are to be reviewed on appeal.

- A. Claims 1-21 stand rejected under 35 U.S.C. § 102(b) as allegedly being anticipated by *Downing et al.* (U.S. 6,289,335).

VII. ARGUMENT

A. Rejection of Claims 1-21 under 35 U.S.C. §102: *Downing et al.*

1. Independent Claim 1

Appellant submits that claim 1 is not anticipated by *Downing et al.* for at least the following reasons. For a proper rejection of a claim under 35 U.S.C. §102, the cited reference

must disclose, teach, or suggest all elements/features/steps of the claim at issue. See, e.g., *E.I. du Pont de Nemours & Co. v. Phillips Petroleum Co.*, 849 F.2d 1430, 7 U.S.P.Q.2d 1129 (Fed. Cir. 1988).

a. ***Downing et al.* does not teach
“logging mechanism configured to maintain a refresh log...containing a first range and a second range that at least partially overlap”**

The final Office Action alleges (p. 2) that this claimed feature is disclosed by *Downing et al.* in Fig. 12, Col. 8 lines 45-63, and Col. 3 line 49 to Col. 4 line 9. Appellant disagrees.

FIG. 12 shows three different logs, none of which shows overlapping ranges (even assuming, *arguendo*, any of these is a “refresh log”). Col. 8 lines 45-63 mentions a master log, but does not mention overlapping ranges (even assuming, *arguendo*, the “master log” is a “refresh log”). The final portion of *Downing et al.* relied on for this teaching does mention various logs:

Another aspect of the invention is a method of refreshing a snapshot. The method includes the step of creating the snapshot based on a plurality of master tables and a snapshot definition query....Each master table has a *primary key, the value of which is recorded into a **master log** in response to detecting a modification to the master table. In response to initiation of a refresh operation, differences between the snapshot and the master tables are reconciled based on the **master log**.*

According to yet another aspect of the invention...In response to detecting a first modification to a first row of the first table, *the primary key value stored in the first row is recorded in a **first log** along with a first value indicative of the first modification. In response to detecting a second modification to a second row of the second table, the primary key value stored in the second row is recorded in a **second log** along with a second value indicative of the second modification. When a refresh operation is initiated, the snapshot is refreshed by reconciling differences between the snapshot, the first table and second table according to the snapshot definition query, the **first log**, the **second log**, the first table, and the second table.*

(*Downing et al.*, Col. 3 line 49 to Col. 4 line 9, emphasis added.)

However, Appellant can find no teaching or even a suggestion that any of these logs contains “a first range and a second range that at least partially overlap” (even assuming, *arguendo*, any of these logs is a “refresh log”).

Thus, *Downing et al.* does not disclose, teach, or suggest all elements of claim 1, and the rejection should be overturned.

b. *Downing et al.* does not teach

“a refresh log...containing a first range and a second range that at least partially overlap...each having a timestamp...respectively indicates when an operation corresponding to the first range and the second range occurred to the table”

The Office Action alleges (p. 3) that this claimed feature is disclosed by *Downing et al.* in Col. 1 first para., Col. 9, last para., and Col. 17 line 45 to Col. 18 line 15. Appellant disagrees.

The first paragraph in Col. 1 states that: “[t]his invention relates to snapshots for database systems and more particularly to a method and system that allows for improved ‘fast refreshes’ of snapshots”. The last paragraph in Col. 9 also deals with refreshes of snapshots, teaching that a refresh may be generated after “a periodic amount of time”. Even assuming that snapshot refreshes involves a refresh log, and that a periodic refresh involves timestamps, these statements do not amount to a teaching that the refresh log contains a timestamp which “indicates when an operation corresponding to the first range and the second range occurred to the table”.

The final portion of *Downing et al.* that the Office Action relies on for teaching this feature does disclose refresh timestamps, as shown below:

In order to prevent master logs from growing indefinitely, entries in the master logs are purged after being used. In a preferred embodiment, however, more than one snapshot may use the same master logs to conserve disk storage for the master logs, so it is important not to purge entries that other snapshots have not yet used. *Accordingly, a refresh timestamp is maintained for each snapshot at the master site. The refresh timestamp for a snapshot indicates the time at which the snapshot was last refreshed.*

Moreover, *each entry in the master logs contains a field for a refresh timestamp, TIME\$\$*. When the entry is first added to a master log, a default value for the timestamp is placed in the TIME\$\$ field. When the entry is first used in a refresh operation, the default value is changed to reflect the time of the refresh operation, as explained hereafter with reference to the flowchart of FIG. 9(a). The default value is preferably a distant time in the future, such as Jan. 1, 4000 A. D. (*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

The purpose of set-up step 800 is to mark the default refresh timestamp values to the current refresh time. With reference to FIG. 9(a), the database system determines the current time (step 900). Then all the master logs for the snapshot are individually processed in a loop (step 902). Each master log is scanned for entries having the default refresh timestamp value in the TIME\$\$\$. In those entries, the value of the TIME\$\$\$ column is set to the current refresh time. When all master logs have been processed, the operation of the set-up step is complete.
(*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

Thus, *Downing et al.* teaches that each entry in a master log is associated with a snapshot refresh timestamp. The refresh timestamp is initialized to a default value. When a master log entry is used in a snapshot refresh, the timestamp is updated to reflect the time of the refresh. The snapshot refresh timestamp in *Downing et al.* thus marks the time that a snapshot is refreshed. Even assuming (for the sake of argument) that refreshing a snapshot involves operations on a table, *Downing et al.* does not teach the specific feature recited in claim 1: “the time stamp associated with each of the first range and second range respectively indicates when an **operation corresponding to the first range and the second range** occurred to the table”. Since *Downing et al.* does not disclose, teach, or suggest all elements of claim 1, the rejection should be overturned.

Furthermore, in the Response to Arguments section (p. 4) the Examiner appears to ignore the plain language of the claim. In the previous response (filed July 9, 2008), Appellant amended to further describe the timestamp associated with each of the first range and the second ranges as indicating “when an operation corresponding to the first range and the second range occurred to the table”. Appellant also explained how this feature further distinguished over *Downing et al.*

Yet in the Response to Arguments section (p. 4) of the final Office Action, the Examiner characterized the Appellant’s argument as “Downing fails to disclose timestamps associated with ranges” – ignoring the amendment. The Response to Arguments then goes on to state:

Examiner finds the limitation range is broad. Examiner interprets the refresh range is the refresh period in which the prior art discloses this feature at col. 1, 1st paragraph; col. 9, last paragraph.
(Office Action, p. 4.)

Regardless of how broadly the Examiner interprets “range”, the Examiner must still give patentable weight to the remainder of the claim limitation which describes the timestamp associated with each of the first range and the second ranges as indicating “when an operation corresponding to the first range and the second range occurred to the table”. As discussed above, *Downing et al.* does not disclose this feature.

2. Independent Claim 5

Applicant notes that the Office Action groups all the independent claims together, such that the same analysis used to reject independent claim 1 is used to reject independent claim 5 – even though a quick perusal of the claims shows they are not coextensive in scope. For example, claim 5 recites “the logging mechanism configured to maintain a refresh log, the refresh log containing a range and a single-row entry, the range and the single row entry each having a timestamp associated therewith”, and this feature is not present in claim 1. Therefore, the Office Action fails to address all the features recited in claim 5. As a result, the rejection of the claims 5-9 is incomplete and deficient, and should be withdrawn.

Nonetheless, in the interest of advancing prosecution, Appellant now discusses distinctions between claim 5 and *Downing et al.* For a proper rejection of a claim under 35 U.S.C. §102, the cited reference must disclose, teach, or suggest all elements/features/steps of the claim at issue. See, e.g., *E.I. du Pont de Nemours & Co. v. Phillips Petroleum Co.*, 849 F.2d 1430, 7 U.S.P.Q.2d 1129 (Fed. Cir. 1988). Claim 5 is not anticipated by *Downing et al.* for at least the reason that *Downing et al.* does not disclose, teach, or suggest a “logging mechanism configured to maintain a refresh log, the refresh log containing a range and a single-row entry, the range and the single row entry each having a timestamp associated therewith, wherein the time stamp associated with the range indicates when an operation corresponding to the range

occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table”.

Downing et al. does disclose refresh timestamps:

In order to prevent master logs from growing indefinitely, entries in the master logs are purged after being used. In a preferred embodiment, however, more than one snapshot may use the same master logs to conserve disk storage for the master logs, so it is important not to purge entries that other snapshots have not yet used. *Accordingly, a **refresh timestamp** is maintained for each snapshot at the master site. The **refresh timestamp** for a snapshot indicates the time at which the snapshot was last refreshed.*

Moreover, *each entry in the master logs contains a field for a **refresh timestamp**, TIME\$\$.* When the entry is first added to a master log, a default value for the timestamp is placed in the TIME\$\$ field. When the entry is first used in a refresh operation, the default value is changed to reflect the time of the refresh operation, as explained hereafter with reference to the flowchart of FIG. 9(a). The default value is preferably a distant time in the future, such as Jan. 1, 4000 A. D. (*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

The purpose of set-up step 800 is to mark the default refresh timestamp values to the current refresh time. With reference to FIG. 9(a), the database system determines the current time (step 900). Then all the master logs for the snapshot are individually processed in a loop (step 902). Each master log is scanned for entries having the default refresh timestamp value in the TIME\$\$ field. In those entries, the value of the TIME\$\$ column is set to the current refresh time. When all master logs have been processed, the operation of the set-up step is complete. (*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

Thus, *Downing et al.* teaches that each entry in a master log is associated with a snapshot refresh timestamp. The refresh timestamp is initialized to a default value. When a master log entry is used in a snapshot refresh, the timestamp is updated to reflect the time of the refresh. The snapshot refresh timestamp in *Downing et al.* thus marks the time that a snapshot is refreshed. Even assuming (for the sake of argument) that refreshing a snapshot involves operations on a table, *Downing et al.* does not teach the specific feature recited in claim 5: “the time stamp associated with the range indicates when an **operation corresponding to the range** occurred to the table and the time stamp associated with the single-row entry indicates when an **operation corresponding to the single-row entry** occurred to the table”. Since

Downing et al. does not disclose, teach, or suggest all elements of claim 5, the rejection should be overturned.

3. Independent Claim 10

For a proper rejection of a claim under 35 U.S.C. §102, the cited reference must disclose, teach, or suggest all elements/features/steps of the claim at issue. *See, e.g., E.I. du Pont de Nemours & Co. v. Phillips Petroleum Co.*, 849 F.2d 1430, 7 U.S.P.Q.2d 1129 (Fed. Cir. 1988). Claim 10 is not anticipated by *Downing et al.* for at least the reason that *Downing et al.* does not disclose, teach, or suggest a “associating a timestamp with the first range and the second range in the refresh log such that the time stamp associated with the first range indicates when an operation corresponding to the first range occurred to the table and the time stamp associated with the second range indicates when an operation corresponding to the second range occurred to the table”. The Office Action alleges (p. 3) that this claimed feature is disclosed by *Downing et al.* in Col. 1 first para., Col. 9, last para., and Col. 17 line 45 to Col. 18 line 15. Appellant disagrees.

The first paragraph in Col. 1 states that: “[t]his invention relates to snapshots for database systems and more particularly to a method and system that allows for improved ‘fast refreshes’ of snapshots”. The last paragraph in Col. 9 also deals with refreshes of snapshots, teaching that a refresh may be generated after “a periodic amount of time”. Even assuming that snapshot refreshes involves a refresh log, and that a periodic refresh involves timestamps, these statements do not amount to a teaching that the refresh log contains a timestamp associated with the first range which “indicates when an operation corresponding to the first range occurred to the table” and a timestamp associated with the second range which “indicates when an operation corresponding to the second range occurred to the table”.

The final portion of *Downing et al.* relied on by the Office Acton as teaching this feature does disclose refresh timestamps:

In order to prevent master logs from growing indefinitely, entries in the master logs are purged after being used. In a preferred embodiment, however, more than one snapshot may use the same master logs to conserve disk storage for the master logs, so it is important not to purge entries that other snapshots have not yet used. *Accordingly, a **refresh timestamp** is maintained for each snapshot at the master site. The **refresh timestamp** for a snapshot indicates the time at which the snapshot was last refreshed.*

Moreover, *each entry in the master logs contains a field for a **refresh timestamp**, TIME\$\$. When the entry is first added to a master log, a default value for the timestamp is placed in the TIME\$\$ field. When the entry is first used in a refresh operation, the default value is changed to reflect the time of the refresh operation, as explained hereafter with reference to the flowchart of FIG. 9(a). The default value is preferably a distant time in the future, such as Jan. 1, 4000 A. D. (Downing et al., Col. 8 lines 45-65, emphasis added.)*

The purpose of set-up step 800 is to mark the default refresh timestamp values to the current refresh time. With reference to FIG. 9(a), the database system determines the current time (step 900). Then all the master logs for the snapshot are individually processed in a loop (step 902). Each master log is scanned for entries having the default refresh timestamp value in the TIME\$\$\$. In those entries, the value of the TIME\$\$\$ column is set to the current refresh time. When all master logs have been processed, the operation of the set-up step is complete. (Downing et al., Col. 8 lines 45-65, emphasis added.)

Thus, *Downing et al.* teaches that each entry in a master log is associated with a snapshot refresh timestamp. The refresh timestamp is initialized to a default value. When a master log entry is used in a snapshot refresh, the timestamp is updated to reflect the time of the refresh.

The snapshot refresh timestamp in *Downing et al.* thus marks the time that a snapshot is refreshed. Even assuming (for the sake of argument) that refreshing a snapshot involves operations on a table, *Downing et al.* does not teach the specific feature recited in claim 10: “associating a timestamp with the first range and the second range in the refresh log such that the time stamp associated with the first range indicates when an operation corresponding to the first range occurred to the table and the time stamp associated with the second range indicates when an operation corresponding to the second range occurred to the table”. Since *Downing et al.* does not disclose, teach, or suggest all elements of claim 10, the rejection should be overturned.

4. Independent Claim 14

Applicant notes that the Office Action groups all the independent claims together, such that the same analysis used to reject independent claim 1 is used to reject independent claim 14 – even though a quick perusal of the claims shows they are not coextensive in scope. For example, claim 14 recites “storing a range and a single-row entry in a refresh log”, and this feature is not present in claim 1. Therefore, the Office Action fails to address all the features recited in claim 14. As a result, the rejection of the claims 14-18 is incomplete and deficient, and should be withdrawn.

Nonetheless, in the interest of advancing prosecution, Appellant now discusses distinctions between claim 14 and *Downing et al.* For a proper rejection of a claim under 35 U.S.C. §102, the cited reference must disclose, teach, or suggest all elements/features/steps of the claim at issue. See, e.g., *E.I. du Pont de Nemours & Co. v. Phillips Petroleum Co.*, 849 F.2d 1430, 7 U.S.P.Q.2d 1129 (Fed. Cir. 1988). Claim 14 is not anticipated by *Downing et al.* for at least the reason that *Downing et al.* does not disclose, teach, or suggest “storing a range and a single-row entry in a refresh log, the range and the single-row entry each having a timestamp associated therewith, wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table occurred to the table”.

The final portion *Downing et al.* that the Office Action relies on as teaching this feature does disclose refresh timestamps:

In order to prevent master logs from growing indefinitely, entries in the master logs are purged after being used. In a preferred embodiment, however, more than one snapshot may use the same master logs to conserve disk storage for the master logs, so it is important not to purge entries that other snapshots have not yet used. *Accordingly, a **refresh timestamp** is maintained for each snapshot at the master site. The **refresh timestamp** for a snapshot indicates the time at which the snapshot was last refreshed.*

Moreover, each entry in the master logs contains a field for a **refresh timestamp**, TIME\$. When the entry is first added to a master log, a default value for the timestamp is placed in the TIME\$ field. When the entry is first used in a refresh operation, the default value is changed to reflect the time of the refresh operation, as explained hereafter with reference to the flowchart of FIG. 9(a). The default value is preferably a distant time in the future, such as Jan. 1, 4000 A. D.
(*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

The purpose of set-up step 800 is to mark the default refresh timestamp values to the current refresh time. With reference to FIG. 9(a), the database system determines the current time (step 900). Then all the master logs for the snapshot are individually processed in a loop (step 902). Each master log is scanned for entries having the default refresh timestamp value in the TIME\$. In those entries, the value of the TIME\$ column is set to the current refresh time. When all master logs have been processed, the operation of the set-up step is complete.
(*Downing et al.*, Col. 8 lines 45-65, emphasis added.)

Thus, *Downing et al.* teaches that each entry in a master log is associated with a snapshot refresh timestamp. The refresh timestamp is initialized to a default value. When a master log entry is used in a snapshot refresh, the timestamp is updated to reflect the time of the refresh. The snapshot refresh timestamp in *Downing et al.* thus marks the time that a snapshot is refreshed. Even assuming (for the sake of argument) that refreshing a snapshot involves operations on a table, *Downing et al.* does not teach the specific feature recited in claim 14: “wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table occurred to the table”. Since *Downing et al.* does not disclose, teach, or suggest all elements of claim 14, the rejection should be overturned.

5. Independent Claim 19

Appellant notes that the functional language in claim 19 does share some similarities to the functions recited in claim 1. Therefore, although the scope of claim 19 is not coextensive with the scope of claim 1, Appellant submits that *Downing et al.* does not disclose, teach, or suggest “the first range and the second range each having a timestamp associated therewith;

wherein the time stamp associated with each of the first range and second range respectively indicates when an operation corresponding to the first range and the second range occurred to the table”, as recited in claim 19, for reasons analogous to those discussed above in connection with claim 1 (section VII.A.1). Since *Downing et al.* does not disclose, teach, or suggest all elements of claim 19, the rejection should be overturned.

6. Dependent Claims 2-4, 6-9, 11-13, 15-18, and 20-21

Since claims 1, 5, 10, 14, and 19 are allowable, Applicants respectfully submit that claims 2-4, 6-9, 11-13, 15-18, and 20-21 are allowable for at least the reason that each depends from an allowable claim. *In re Fine*, 837 F.2d 1071, 5 U.S.P.Q. 2d 1596, 1598 (Fed. Cir. 1988). Therefore, Applicants respectfully request that the rejection of claims 2-4, 6-9, 11-13, 15-18, and 20-21 be withdrawn.

Dependent claims 2-3, 6-8, 11-12, 15-17, and 20-21 are allowable for the separate and independent reason that *Downing et al.* does not disclose, teach, or suggest a “refresh log comprising an epoch identifier”. In the Response to Arguments section (p. 5), the Examiner notes that para. 32 of the instant specification discloses “that the epoch number may be used to identify a group of rows that have been added to the IUD log since a refresh operation was performed”, then refers to Figs. 13a-c in *Downing et al.* as disclosing “a series of states of a fast refreshed snapshot”.

Appellant does not understand exactly what the Examiner believes Figs. 13a-c to teach. Appellant assumes (for the sake of argument) that Figs. 13a-c show rows being added to a refresh log. Even so, the action of adding rows to a refresh log is not the same the log including a particular identifier (as recited in claims 2-3, 6-8, 11-12, 15-17, and 20-21). Appellant further submits that none of the identifiers in Figs. 13a-c would be understood by a person of ordinary skill in the art as an “epoch identifier”. To the extent that the Response to Arguments section actually alleges that one of the identifiers in Figs. 13a-c is the same as an epoch identifier,

Appellant submits that the Examiner has merely made a conclusory statement that the identifiers in *Downing et al.* are the same as “epoch identifiers”, rather than providing any reasoning or evidentiary foundation as to why a person of ordinary skill in the art would understand these to be the same.

B. Conclusion

For at least the reasons discussed above, Appellant respectfully requests that the Examiner's final rejection of claims 1-21 be overturned by the Board. In addition to the claims listed in Section VIII (CLAIMS – APPENDIX), Section IX (EVIDENCE – APPENDIX) included herein indicates that there is no additional evidence relied upon by this brief. Section X (RELATED PROCEEDINGS – APPENDIX) included herein indicates that there are no related proceedings.

Respectfully submitted,

By: /Karen G. Hazzah/

Karen G. Hazzah,
Reg. No. 48,472

**THOMAS, KAYDEN, HORSTEMEYER
& RISLEY, L.L.P.**

600 Galleria Parkway, SE
Suite 1500
Atlanta, Georgia 30339-5948
Tel: (770) 933-9500
Fax: (770) 951-0933

VIII. CLAIMS – APPENDIX

1. A system, comprising:
 - a processor;
 - a memory;
 - a materialized view stored on the memory, the materialized view being derived at least in part from a table;
 - a logging mechanism stored on the memory, the logging mechanism configured to maintain a refresh log, the refresh log containing a first range and a second range that at least partially overlap, the first range and the second range each having a timestamp associated therewith; wherein the time stamp associated with each of the first range and second range respectively indicates when an operation corresponding to the first range and the second range occurred to the table; and
 - a refresh manager stored on the memory, the refresh manager configured to resolve conflicts between the first range and the second range that at least partially overlap by selecting portions of the first range and the second range that have the more recent timestamp and applying the selected portions of the first range and the second range to the materialized view.
2. The system set forth in claim 1, wherein the refresh log comprises a plurality of entries, each of the entries comprising an epoch identifier.
3. The system set forth in claim 2, wherein the epoch identifier is defined to correspond to changes that have been made to the table since a previous refresh operation on the materialized view.

4. The system set forth in claim 1, wherein a plurality of materialized views are derived at least in part from the table.

5. A system, comprising:

a processor;

a memory;

a materialized view stored on the memory, the materialized view being derived at least in part from a table;

a logging mechanism stored on the memory, the logging mechanism configured to maintain a refresh log, the refresh log containing a range and a single-row entry, the range and the single row entry each having a timestamp associated therewith, wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table and;

a refresh manager stored on the memory, the refresh manager configured to resolve conflicts between the range and the single-row entry by ignoring the single-row entry if the single-row entry is part of the range and if the single-row entry has the more recent timestamp and by applying the single-row entry to the materialized view if the single-row entry is not part of the range or if the range has the more recent timestamp.

6. The system set forth in claim 5, wherein the refresh log comprises a plurality of entries, each of the entries comprising an epoch identifier.

7. The system set forth in claim 6, wherein the epoch identifier is defined to correspond to changes that have been made to the table since a previous refresh operation on the materialized view.
8. The system set forth in claim 7, wherein the single-row record belongs to an epoch E, a latest screening range belongs to an epoch $E' < E$, and the refresh manager is adapted to ignore the single-row record for a materialized view that fulfils $MV.EPOCH[T] \leq E'$ and to apply the single-row record to a materialized view that fulfils $MV.EPOCH[T] > E'$.
9. The system set forth in claim 5, wherein a plurality of materialized views are derived at least in part from the table.
10. A method, comprising:
 - deriving a materialized view at least in part from a table;
 - storing a first range and a second range that at least partially overlap in a refresh log;
 - associating a timestamp with the first range and the second range in the refresh log such that the time stamp associated with the first range indicates when an operation corresponding to the first range occurred to the table and the time stamp associated with the second range indicates when an operation corresponding to the second range occurred to the table; and
 - resolving conflicts between the first range and the second range by applying a portion of either the first range or the second range that has the more recent timestamp to the materialized view.
11. The method for performing conflict resolution set forth in claim 10, comprising creating a plurality of records in the refresh log and storing an epoch identifier in each of the records.

12. The method for performing conflict resolution set forth in claim 11, comprising defining the epoch identifier to correspond to changes that have been made to the table since a previous refresh operation on the table.
13. The method for performing conflict resolution set forth in claim 10, comprising deriving a plurality of materialized views at least in part from the table.
14. A method, comprising:
- deriving a materialized view at least in part from a table;
 - storing a range and a single-row entry in a refresh log, the range and the single-row entry each having a timestamp associated therewith, wherein the time stamp associated with the range indicates when an operation corresponding to the range occurred to the table and the time stamp associated with the single-row entry indicates when an operation corresponding to the single-row entry occurred to the table;
 - ignoring the single-row entry if the single-row entry is part of the range and if the single-row entry has the more recent timestamp; and
 - applying the single-row entry to the materialized view if the single-row entry is not part of the range or if the range has the more recent timestamp.
15. The method set forth in claim 14, comprising storing a plurality of entries in the refresh log, each of the plurality of entries comprising an epoch identifier.
16. The method set forth in claim 15, comprising defining the epoch identifier to correspond to changes that have been made to the table since a previous refresh operation on the materialized view.

17. The method set forth in claim 16, wherein the single-row record belongs to an epoch E, a latest screening range belongs to an epoch $E' < E$, the method comprising: ignoring the single-row record for a materialized view that fulfils $MV.EPOCH[T] \leq E'$; and applying the single-row record to a materialized view that fulfils $MV.EPOCH[T] > E'$.

18. The method set forth in claim 14, comprising deriving a plurality of materialized views at least in part from the table.

19. A computer program, comprising:

a machine readable medium; a logging mechanism stored on the machine readable medium, the logging mechanism being adapted to create a refresh log that contains a first range and a second range that at least partially overlap, the first range and the second range each having a timestamp associated therewith wherein the time stamp associated with each of the first range and second range respectively indicates when an operation corresponding to the first range and the second range occurred to the table; and

a refresh manager stored on the machine readable medium, the refresh manager being adapted to resolve conflicts between the first range and the second range that at least partially overlap by selecting portions of the first range and the second range that have the more recent timestamp and applying the selected portions of the first range and the second range to the materialized view.

20. The computer program set forth in claim 19, wherein the refresh log comprises a plurality of entries, each of the entries comprising an epoch identifier.

21. The computer program set forth in claim 20, wherein the epoch identifier is defined to correspond to changes that have been made to the table since a previous refresh operation on any materialized view that is derived at least in part from the table.

IX. EVIDENCE – APPENDIX

None.

X. RELATED PROCEEDINGS – APPENDIX

None.